

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

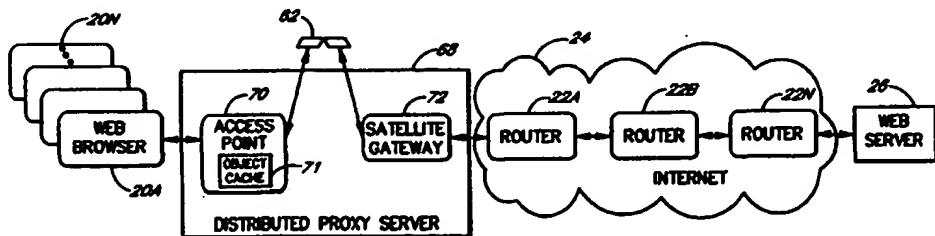
**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :	A1	(11) International Publication Number: WO 99/08429
H04L 29/06		(43) International Publication Date: 18 February 1999 (18.02.99)
(21) International Application Number: PCT/US98/16353		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 5 August 1998 (05.08.98)		
(30) Priority Data: 60/054,847 6 August 1997 (06.08.97) US		
(71) Applicant: TACHYON, INC. [US/US]; Suite 101, 6225 Nancy Ridge Drive, San Diego, CA 92121 (US).		
(72) Inventors: CARNEAL, Bruce, L.; 13172 Caminito Pointe Del Mar, Del Mar, CA 92014 (US). FISH, Ram, A.; 820 West "G" Street #414, San Diego, CA 92101 (US). GURALNICK, Jeremy; 832 Sunrich Lane, Encinitas, CA 92024 (US).		Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>
(74) Agent: ALTMAN, Daniel, E.; Knobbe, Martens, Olson and Bear, LLP, 16th floor, 620 Newport Center Drive, Newport Beach, CA 92660 (US).		

(54) Title: A DISTRIBUTED SYSTEM AND METHOD FOR PREFETCHING OBJECTS



(57) Abstract

In an internet access system which includes a satellite link, a distributed proxy server (68) is provided which reduces a delay associated with the retrieval of inline objects of web pages. The distributed proxy server (68) includes an access point component (70) and a satellite gateway component (72). The access point component (70) runs on the client (browser) side of the satellite link and communicates with web browsers (20A-20N). The satellite gateway component (72) runs on the internet side of the satellite link and communicates with web servers (26). As a web page is retrieved over the satellite link, the satellite gateway component (72) parses the base file component of the web page to identify any references to inline objects of the web page, and prefetches each such inline object. The distributed proxy server thereby eliminates the delays normally associated with (a) waiting for the web browser (20) to receive the base component and request the object(s), and (b) waiting for the browser's object request(s) to be transmitted over the satellite link. The prefetched objects are transmitted over the satellite link to the access point component (70), which in-turn stores the prefetched objects in an object cache (71). When a web browser (20) requests an inline object, the access point component (70) checks the cache (71), and if the object resides therein, returns the object to the browser (20) without forwarding the object request over the satellite link. Traffic over the forward satellite link is thus reduced. The method implemented by the distributed proxy server (68) can also be used to reduce delays and traffic over other types of links, including non-wireless links.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon			PT	Portugal		
CN	China	KR	Republic of Korea	RO	Romania		
CU	Cuba	KZ	Kazakhstan	RU	Russian Federation		
CZ	Czech Republic	LC	Saint Lucia	SD	Sudan		
DE	Germany	LI	Liechtenstein	SE	Sweden		
DK	Denmark	LK	Sri Lanka	SG	Singapore		
EE	Estonia	LR	Liberia				

**A DISTRIBUTED SYSTEM AND METHOD FOR
PREFETCHING OBJECTS**

5

10

Background of the Invention

I. Field of the Invention

This invention relates generally to data transfer. More specifically, the invention relates to digital data transfer over a digital network.

15 **II. Description of the Related Art**

The bloom of the Internet has encouraged many companies and individuals to establish an Internet presence. For example, a company may create a web page which describes its products and services and allows a user to place a purchase order. These web pages are stored on web servers. A user may access a web page from the a web server using web browser software running on a computer. The web page may contain links to other information at the same site or other web sites.

20 Figure 1 is a block diagram showing an Internet connection. A user originates a file request from a web browser 20. The web browser 20 may comprise a personal computer, a network terminal or any other manner of digital user terminal capable of executing web browsing software. The request is passed through a series of routers 22A-22N of the Internet 24. The routers 22A-22N do not examine the contents of the request but simply transfer the request to an appropriate web server 26 according to an address header. The web server 26 examines the contents of the request and responds with the requested file.

25 When a user would like to access information on the Internet, the user enters a uniform resource locator (URL) into the web browser. The URL is basically a pointer

to the location of an object. For example, "http:\\www.internic.net\\rfc\\rfc1738.txt" is the URL address which points to a Request For Comment document which describes uniform resource locators. In the URL, the "http" indicates that the HyperText Transfer Protocol (HTTP) protocol is used to access the site. A double backslash indicates that a host name follows such as "www.internic.net". A single backslash indicates that either a directory or a filename follows. In this case, "rfc" is a directory and "rfc1738.txt" is the file in that directory which is displayed when this URL is requested by the web browser 20.

The World Wide Web is built on top of the Internet. HTTP is the client/server protocol used most commonly on the World Wide Web. HTTP is used to set up communication between a client and a server and pass commands and files between the two systems.

HTTP provides a means for a web browser to access a web server and request documents created using the HyperText Markup Language (HTML). HTML web pages can include images, sound clips, text files and other types of objects. Some of the objects may not be part of the original HTML parent file (the base component of the Web document) requested by the web browser 20. Instead, the HTML parent file contains external references to these inline objects, which are in the form of other data files on the server. When a user retrieves the HTML parent file on the web browser, the inline objects are also retrieved and inserted into the displayed of the document. Thus, an HTML document (or "page") actually consists of the HTML parent file along with any additional sound, graphics and multimedia inline objects specified with the parent file. For example, the inline objects may include advertising banners, sliders, bullet listings, graphic images, sound clips or other such items.

Figure 2 is a timing diagram showing data transfer to and from the web browser 20. In Figure 2, time progresses from left to right. The upward pointing arrows indicate outgoing messages from the web browser 20 intended for the web server 26. Downward arrows indicate incoming messages received at the web browser 20 from the web server 26. For simplicity of illustration, each incoming and outgoing message appears to be transferred instantaneously. In actual implementations, the transfer of each message typically requires a discernible amount of time.

An outgoing message 30 carries the initial URL request. In response, an incoming message 32 carries the first portion of a response to the request carried in the outgoing message 30. An incoming message 34 and an incoming message 36 correspond to a second and third portion of the response.

5 Assume that the incoming message 32 contains an external reference to an inline object. The web browser 20 examines the incoming information and in response sends an outgoing message 38 which carries a request for the inline object. For illustration purposes, we shall assume that the inline object is a sound clip.

10 Following the outgoing message 38, the web browser 20 receives an incoming message 40 containing additional information corresponding to the initial request carried in the outgoing message 30. After reception of the incoming messages corresponding to the initial request, the web browser 20 begins to receive the sound clip within an incoming message 42. In an incoming message 44, the web browser 20 continues to receive information concerning the sound clip.

15 Assume that the incoming message 42 contains an external reference to an inline object which is an ad banner. An outgoing message 46 carries a request for the ad banner. Following the outgoing message 46, the web browser 20 receives an incoming message 48 and an incoming message 50 containing additional information corresponding to the sound clip. Finally, in an incoming message 52, the web browser 20 receives the information concerning the ad banner.

20 Each time that the web browser 20 requests information from the web server 26, a delay is incurred. For example, notice that a time delay ΔT_{1A} elapses between the outgoing message 30 and the corresponding incoming message 32. The delay includes two primary components: (i) the round-trip delay associated with connection to the web server 26 and (ii) the response time of the web server 26. In the Figure 2 example, the transfers of the inline objects are delayed by transfers of previously requested objects and the parent file and the time delays ΔT_{2A} and ΔT_{3A} are, therefore, longer than the delay ΔT_{1A} .

25 As described in more detail below, because the HTTP protocol requires the web browser to examine the parent file and generate separate requests for the inline objects, the introduction of a link which introduces significant delay can greatly increase the

5

amount of time required to fully retrieve and display a web page. For example, if the user's internet access channel includes a satellite link, the time required to retrieve a web page that includes a single inline object will be at least twice the round-trip delay of the satellite link. Further, the need to separately request inline objects produces unnecessary traffic over the communications link. The present invention seeks to overcome these problems without the need to modify the HTTP protocol.

Summary of the Invention

The present invention addresses the above problems by providing a distributed system and method for prefetching inline objects of documents. In a preferred embodiment, the system is in the form of a distributed proxy server for use in an internet access system which includes a satellite link. The distributed proxy server includes an access point component which runs on the client (browser) side of the satellite link and communicates with web browsers, and includes a satellite gateway component which runs on the internet (web server) side of the satellite link and communicates with web servers. In operation, when a web server returns a parent file of a web page that has been requested by the user, the satellite gateway component parses the parent file to identify any references to inline objects, and prefetches these objects from the web server. The objects are thus requested without waiting for the browser to receive the parent file and generate requests for the inline objects.

20

The satellite gateway forwards the prefetched objects over the satellite link to the access-point component, which in-turn caches the inline objects until requested by the browser. If the access point component receives a request for an object which resides in the cache, the access point component returns the object without allowing the object request to be transmitted over the satellite link. The distributed proxy server thus reduces the delay associated with requests for inline objects, and reduces traffic over the satellite link.

25

Although the system in the preferred embodiment operates in conjunction with a satellite link, the underlying method and architecture can also be used to increase performance over other types of links, including non-wireless links. In addition, although the preferred embodiment operates in a system which uses HTTP, the invention

can also be used with other types of document retrieval protocols in which inline objects are requested separately from the base component.

In accordance with the invention, there is thus provided, in a client-server type document retrieval system in which inline objects of documents are requested and retrieved separately from base components of the documents, a distributed system for reducing a performance degradation caused by a communications link. The distributed system comprises a first component which runs on the client side of the communications link and communicates with clients, the first component being adapted to receive document requests from the clients and to forward the requests over the communications link for processing. The system also includes a second component which runs on the server side of the communications link and communicates with document servers, the second component being adapted to receive the document requests from the first component over the communications link and to forward the requests to the document servers, the requests causing the document servers to return base components of requested documents. In operation, the second component processes base components returned by the document servers by at least (i) parsing the base components to identify references to inline objects, (ii) prefetching the inline objects, and (iii) forwarding the base components and prefetched inline objects to the first component. The first component stores the prefetched inline objects received from the second component in a cache memory, and responds to object requests from the clients by forwarding the inline objects to the clients from the cache memory.

Brief Description of the Drawings

The features, objectives, and advantages of the invention will become more apparent from the detailed description set forth below when taken in conjunction with the drawings wherein like parts are identified with like reference numeral throughout and wherein:

Figure 1 is a block diagram showing an Internet connection.

Figure 2 is a timing diagram showing data transfer to and from a web browser.

Figure 3 is a block diagram showing an Internet connection comprising a satellite link.

Figure 4 is a timing diagram illustrating the transmission of the data shown in Figure 2 after introduction of a satellite link into the system.

5 Figure 5 is a block diagram showing an Internet connection comprising a proxy server.

Figure 6 is a block diagram showing a digital data network comprising a distributed proxy server according to the invention.

10 Figure 7 is a data flow diagram showing one embodiment of data transfer according to the invention.

Figure 8 is a timing diagram illustrating the transmission of the data shown in Figures 2 and 4 as transferred over a distributed proxy server according to the invention.

Detailed Description of a Preferred Embodiment

The use of wireless communication systems for the transmission of digital data is becoming more and more pervasive. In a wireless system, the most precious resource in terms of cost and availability is typically the wireless link itself. Therefore, one major design goal in designing a communication system comprising a wireless link is to efficiently use the available capacity of the wireless link. In addition, typically, the delay associated with traversing the wireless link is significantly larger than the delay associated with the remainder of the network. Therefore, it is also desirable to reduce the delay associated with use of a wireless link.

20 Although stand-alone computers are very powerful tools, they become even more powerful when coupled together to form a network. As an increasing number of computers are coupled together, increasing demands are placed on the network which couples them together. If the computers are located close to one another, they may be interconnected by dedicated wireline connections. A computer which is located some significant distance away from a network may access the network by connection through a standard telephone line. However, telephone lines have a limited bandwidth which places a limit on the rate at which data may be transferred between the computer and the network.

25 Therefore, alternative means of accessing digital data networks have been developed. For example, satellite links may be used to transfer digital data within a data

5

network. The use of a geosynchronous satellite link introduces a round-trip delay approximately equal to one-half of a second. The protocols currently in use on typical digital data networks were not developed with such large round-trip delays in mind. The use of a link which introduces significant delay can decrease the average data transfer rate below an acceptable level. The present invention reduces the inefficiencies caused by using a link which introduces significant delay to access a digital data system. The present invention also increases the efficient use of the digital data system link.

10

15

Figure 3 is a block diagram illustrating an Internet connection comprising a satellite link. In Figure 3, the web browser 20 is coupled to a satellite terminal 60. The satellite terminal 60 is coupled to an earth station 64 via a geosynchronous satellite 62. The earth station 64 is coupled to the series of routers 22A-22N of the Internet 24. The Internet 24 is coupled to the web server 26. Inclusion of the geosynchronous satellite link introduces a delay of about 270 milliseconds per transmission of a signal between the satellite terminal 60 and the earth station 64. Thus, each data transfer between the web browser 20 and the web server 26 exhibits a round-trip delay of at least one-half of a second.

20

Figure 4 is a timing diagram illustrating the transmission of the data shown in Figure 2 over a system comprising a satellite link. Due to the nature of the nested requests for the inline objects, not only does the system incur an initial system delay equal to the round-trip delay, but it also incurs additional delay as the subsequent requests are made for the inline objects.

25

30

An incoming message 30' carries the initial URL request which is forwarded over the wireless link. Due to the round trip delay of the link, an incoming message 32' which carries the first portion of a response to the request carried in the outgoing message 30' is received after a delay of time delay ΔT_{1B} . Assume that an incoming message 32' comprises an external reference to an inline object which is a sound clip. In response, the web browser 20 sends an outgoing message 38' which carries a request for the inline object. Due to the round trip delay of the link, an incoming message 42' which carries the first portion of a response to the request carried in the outgoing message 38' is received after a delay of time delay ΔT_{2B} . Assume that a responsive incoming message 42' contains an external reference to an inline object which is an ad

banner. An outgoing message 46' carries a request for the ad banner. Due to the round trip delay of the link, an incoming message 52' which carries the first portion of a response to the request carried in the outgoing message 46' is received after a delay of time delay ΔT_{3B} .

5 Notice that the time delays ΔT_{1B} , ΔT_{2B} , and ΔT_{3B} are significantly longer in Figure 4 than the time delays ΔT_{1A} , ΔT_{2A} , and ΔT_{3A} in Figure 2. This difference is due chiefly to the round trip delay associated with the wireless link. Each time that a request is made, the first response is delayed by at least the round trip delay. Thus, the inclusion of a link which introduces a relatively long delay can significantly decrease
10 the average data transfer rate as well as introduce a significant initial fixed delay.

15 Figure 5 is a block diagram showing an Internet connection comprising a proxy server 66. The proxy server 66 is designed to interface a plurality of web browsers 20A-20N to the Internet 24. The proxy server 66 may be designed to protect the web browsers 20A-20N (and the network upon which they run) from hostile invasion by way
20 of the Internet 24. For example, the proxy server 66 may be a one-way server that blocks other Internet users from accessing the internal network. Only packets that are received in response to an internal user's request are allowed back through the proxy server 66 from the Internet 24. Other data received at the proxy server 66 is not passed on to the web browsers 20A-20N.

25 Because the proxy server 66 provides a single point of connection for a set of users, it is fairly easy to perform functions such as virus scanning, content filtering and access control at the proxy server 66 rather than at the individual web browsers 20A-20N. Thus, a company may retain control over its internal network more easily using the proxy server 66 as a central control point for the network.

20 The proxy server 66 is designed to be transparent to the system users meaning that the web browsers 20A-20N are not aware of the existence of the proxy server 66. Web access occurs seamlessly through the proxy server 66. Although the web browsers 20A-20N may be unaware of the proxy server 66, the end result is that the web browsers 20A-20N receive a combination of information from the proxy server 66 and the
30 Internet 24.

Often times, caching is used to increase the performance of the system as perceived by the web browser 20. The term "caching" refers to the process of storing copies of documents received by the proxy server 66 onto a local storage media (typically a disk but also main memory for short term caching.) Because the proxy server 66 provides a central location where each of the web browsers 20A-20N accesses the Internet 24, the proxy server 66 may cache frequently accessed documents so that they are quickly available to other web browsers 20 which request them. For example, in a financial services organization, many different browsers may access the Wall Street Journal home page during any given day. If the proxy server 66 has caching ability, the Wall Street Journal home page may be cached after the first request of each day. Subsequent requesters receive the Wall Street Journal home page from the local cache rather than over the Internet 24. In general, caching improves the performance of the system by reducing latency and saving network bandwidth.

One kind of caching, called "prefetching", involves caching documents that are likely to be requested by a client. One example of a document that is likely to be requested by a client is an inline object within a previously requested parent file. If the proxy server parses the originally parent file, the proxy server can determine that the document contains an external reference to an inline object and can begin retrieval of the inline object before the actual request arrives from the web browser.

In the case of prefetching inline objects, the request prediction is highly deterministic. In the rare case where the web browser software is not configured to load inline objects automatically, the prediction will fail. Also, if the user interrupts the retrieval before all inline objects are requested, some of the objects will not be requested by the web browser. However, in the general case, each of the inline objects referenced by the parent document is requested by the web browser.

If one attempts to integrate a prior art proxy server comprising a caching ability into an Internet connection comprising a wireless satellite link, system performance is not greatly improved. For example, referring again to Figure 3, assume that a proxy server is inserted between the earth station 64 and the router 22A. When an initial request is generated by the web browser 20, it is forwarded from the satellite terminal 60 over the wireless link to the earth station 64. The earth station 64 forwards the

request to the proxy server which in turns forwards it to the Internet 24. When the response from the Internet 24 is received, the proxy server forwards it to the earth station 64. The proxy server also parses the information in the response to the initial request to determine the presence of inline objects. The proxy server sends surrogate requests for the inline objects and stores the objects as they arrive. Meanwhile, the response to the initial request is forwarded over the wireless link to the web browser 20. When the web browser 20 receives the response, it forwards a request for the inline object back over the wireless link. When the proxy server receives the request, it responds with the cached prefetched information, thus, eliminating any extra delay associated with transmission over the Internet 24. The information corresponding to the inline object traverses the wireless link to the web browser 20. Thus, the round trip delay of the wireless link is not avoided. Because the round-trip delay over the wireless link is the most significant factor in determining the latency of the system in Figure 3, system performance is not significantly improved by the inclusion of a proxy server on the gateway side of the satellite link.

Likewise, assume that a prior art proxy server is inserted between the satellite terminal 60 and the web browser 20. When an initial request is generated by the web browser 20, it is forwarded to the proxy server which forwards it to the satellite terminal 60. The satellite terminal 60 forwards it over the wireless link to the earth station 64. The earth station 64 forwards the request to the Internet 24. When the response from the Internet 24 is received, the earth station 64 forwards the response over the wireless link to the satellite terminal 60 and the proxy server. The proxy server forwards the response to the web browser 20. The proxy server also parses the information to determine the presence of inline objects. The proxy server prefetches the inline objects by sending a surrogate request for them over the wireless link. When the web browser 20 receives the initial response, it forwards a request for the inline object to the proxy server. When the proxy server receives the request from the web browser 20, it intercepts the request and does not forward it over the wireless link. The proxy server waits for the surrogate request and the corresponding response to traverse the wireless link and forwards the information to the web browser when available, thus, eliminating any extra delay associated with transmission of the initial response from the proxy

5

server to the web browser 20 and transmission of the request from the web browser 20 to the proxy server. However, the round trip delay of the wireless link is not avoided. Because the round-trip delay over the wireless link is the most significant factor in determining the latency of the system in Figure 3, system performance is not significantly improved by the inclusion of a proxy server on the satellite terminal side of the satellite link.

Figure 6 is a block diagram showing a digital data network comprising a distributed proxy server according to the invention. In Figure 6, the web browsers 20A-20N are coupled to the Internet 24 by means of a distributed proxy server 68. The distributed proxy server 68 comprises two components: an access point component 70 ("access point") and a satellite gateway component 72 ("satellite gateway"). The access point 70 and the satellite gateway 72 may be implemented on one or more workstations or personal computers which run corresponding software components. Some or all of the functions of the access point 70 and satellite gateway 72 can alternatively be implemented using application-specific hardware. The access point 70 and the satellite gateway 72 are connected via a wireless link over the satellite 62. The web browsers 20A-20N are preferably connected to the access point 70 by a local area network (not shown). As depicted in Figure 6, the access point 70 implements an object cache 71 for temporarily storing objects that have been prefetched by the satellite gateway 72.

10

The distributed proxy server 68 facilitates efficient data transfer by reducing the effect of the round-trip delay associated with transversal of the satellite link. The access point 70 and the satellite gateway 72 may also provide services typically performed by prior art proxy servers. In the preferred embodiment, the use of the distributed proxy server 68 is transparent to the entities coupled to the Internet 24, such as the web server 26 and the web browser 20.

15

Figure 7 is a data flow diagram showing one embodiment of data transfer according to the invention in which the requested web page is comprised of a parent file component which references an inline object. In block 80, the web browser 20 requests a web page. The access point 70 receives the request and forwards it to the satellite gateway 72 over the wireless link in block 82. After some delay, the request is received at the satellite gateway 72. In block 84, the satellite gateway 72 forwards the request to

20

25

the Internet 24. The satellite gateway 72 may comprise a network communication module for executing this function.

The web server 26 receives the request and responds with the parent file. The satellite gateway 72 receives the parent file from the Internet 24 and forwards it to the access point 70 in block 86. After the file traverses the wireless link, the access point 70 forwards the parent file to the web browser 20 in block 88. In block 90, the web browser receives the parent file.

Prior to, during or after transmission of the transfer of block 86, the satellite gateway parses the parent file, extracts an external reference to an inline object and forwards a surrogate request for the inline object over the Internet 24 in block 92. For example, the satellite gateway 72 may comprise a parsing module and a surrogate request generation module to carry out these functions. Also in block 92, the satellite gateway 72 sends a message to the access point 70 identifying the inline object which it is prefetching. For example the satellite gateway 72 may comprise a message module which creates and sends messages to the access point 70. In block 94, the access point 70 receives the prefetch object listing and stores it for later reference, such as within a caching module.

In block 96, the web browser 20 forwards a request for the inline object. In block 98, the access point 70 compares the request to the list of prefetch objects. If the request corresponds to an entry on the list of prefetch objects, the request is intercepted and is not forwarded over the wireless link. If the requested object does not correspond to an entry on the list and is not available in the cache, the access point 70 forwards the request over the wireless link. This function may be carried out by a message analysis and comparison module.

In the Figure 7 example, because the inline object appears on the list of objects which have been or are currently being fetched, the access point 70 intercepts the request and does not forward the request over the satellite link. Meanwhile, in block 100, the inline object is received at the satellite gateway 72 and automatically forwarded to the access point 70. Notice that the satellite gateway 72 has not received a request for the inline object and is sending a response to the initial request to the access point 70 containing information which the access point 70 has not requested from the satellite

gateway 72. Such a response is termed a "chubby response." For example, the satellite gateway 72 may comprise a chubby response creation module which performs these functions. In block 102, the access point 70 receives the object and forwards it to the web browser 20. In block 104, the web browser 20 receives the inline object. In this way, the delay associated with awaiting the transmission of the request over the satellite link is avoided and the inline objects are transferred shortly after the parent file. In addition, the bandwidth normally occupied by the wireless requests for inline objects is freed for other uses.

The efficiencies gained by using data flow as described in Figure 7 can be readily seen with reference to the timing diagram of Figure 8. Figure 8 is a timing diagram illustrating the transmission of the data shown in Figures 2 and 4 as transferred over a distributed proxy server according to the invention. In order to illustrate the data flow more clearly, in addition to the data flow as perceived at the web browser 20, Figure 8 also shows the data flow as perceived by the access point 70 and by the satellite gateway 72.

An outgoing message 130 from the web browser carries the initial URL request. The access point 70 forwards a corresponding outgoing message 130' over the wireless link. The satellite gateway 72 sends a corresponding outgoing message 130'' over the Internet 24. Notice the delay between the outgoing message 130' and the outgoing message 130'' due to the introduction of the wireless link. An incoming message 132'' carries the first portion of a response to the request of the outgoing message 130''. An incoming message 134'' and an incoming message 136'' carry a second and third portion of the response. The incoming message 132'' carries an external reference to an inline object which we shall assume is a sound clip. The satellite gateway 72 parses the incoming message 132'' and recognizes the external reference. In an outgoing message 138'', the satellite gateway 72 sends a surrogate request over the Internet 24 for the sound clip, in this way prefetching the sound clip.

The incoming messages 132'', 134'' and 136'' are forwarded to the access point 70 where they are received as the incoming messages 132', 134' and 136', respectively, once again incurring a delay due to the wireless link. The access point 70 forwards them to the web browser 20 where they are received as the incoming messages 132, 134

and 136, respectively. In response to the incoming message 132, the web browser 20 sends an outgoing message 138 requesting the sound clip. Notice that at the time the outgoing message 138 is sent from the web browser 20, the outgoing message 138' has already been forwarded by the satellite gateway 72 to the Internet 24. Therefore, the 5 access point 70 intercepts the request and does not forward it to the gateway 72, thus, reducing the amount of data sent over the satellite link. Therefore, in Figure 8, no corresponding outgoing message 138' is forwarded from the access point 70 to the satellite gateway 72. At the satellite gateway 72, an incoming message 140'' is received as the fourth and last response to the original request and is forwarded to the access 10 point 70 where it is received as an incoming message 140'.

An incoming message 142'' and an incoming message 144'' are received in response to the request for the sound clip. The incoming message 142'' carries an external reference to an inline object which we shall assume is an ad banner. Once again, the satellite gateway 72 parses the incoming message 142'' and detects external 15 reference to the ad banner. An outgoing message 146'' is sent by the satellite gateway 72 carrying a surrogate request for the ad banner.

Meanwhile, the incoming messages 142'' and 144'' are automatically forwarded to the access point 70 without awaiting a corresponding request from the access point 70 for the inline object. In this way, the inline objects are both prefetched and pre-transferred over the wireless link, thus, significantly reducing the latency perceived by 20 the web browser 20. The incoming messages 142'' and 144'' are received as the incoming messages 142' and 144', respectfully, at the access point 70. The incoming messages 142' and 144' are forwarded to the web browser 20 after they are requested by the web browser 20 in the outgoing message 138. They are received by the web 25 browser 20 as the incoming messages 142 and 144, respectfully. In response to receipt of the incoming message 142 containing the external reference to the ad banner, the web browser 20 sends an outgoing message 146 requesting the ad banner. Once again, the access point 70 intercepts the request and no corresponding outgoing message 146' is forwarded from the access point 70 to the satellite gateway 72.

30 Meanwhile, the satellite gateway 72 continues to receive an incoming message 148'' and an incoming message 150'' in response to the request for the sound clip and

forwards them to the access point 70. The access point 70 receives them as the incoming messages 148' and 150' subject to the delay introduced by the wireless link. The access point 70 forwards these transmissions to the web browser 20 where they are received as the incoming messages 148 and 150. In response to the request in the 5 outgoing message 146" for the ad banner, an incoming message 152" is received at the satellite gateway 72 and is automatically forwarded to the access point 70 without awaiting a corresponding request. The access point 70 receives the ad banner information as an incoming message 152' and forwards it to the web browser 20 in response to a request carried in the outgoing message 146 where it is received as an 10 incoming message 152.

Examination of Figure 8 in comparison with the prior art of Figure 4 reveals the efficiencies introduced by the use of the present invention. Notice that there is a substantial delay ΔT_{1C} between the incoming message 130 and the incoming message 132 in Figure 8 just as there is a substantial delay ΔT_{1A} between the incoming message 30' and the incoming message 32' in Figure 4. A large portion of this delay is due to the round-trip delay associated with using the wireless link. However, note that the substantial delay ΔT_{2B} between the outgoing message 38' and the incoming message 42' of Figure 4 is not exhibited by the delay ΔT_{2C} between the outgoing message 138 and the incoming message 142 in Figure 8 due to the fact that the information was prefetched and also pre-transferred in a chubby response. In addition, the substantial delay ΔT_{3B} between the outgoing message 46' and the incoming message 52' of Figure 4 is not exhibited by the delay ΔT_{3C} between the outgoing message 146 and the incoming message 152 of Figure 8 for the same reason. The delays ΔT_{2C} and ΔT_{3C} in Figure 8 are comparable to the delays ΔT_{2A} and ΔT_{3A} in Figure 2. In this way, the average data rate is 15 increased by use of the invention and the latency associated with system use is 20 decreased. 25

In Figure 7, the access point 70 has been notified that the inline object has been prefetched in block 94 before the request for the inline object is made by the web browser 20 in block 96. However, in some cases, the web browser 20 may request the 30 inline object before the prefetch notification is received. In such a case, the access point 70 may store the request for some preprogrammed period of time such as three seconds.

If within the period of time the object does not appear on a list of prefetch objects nor is the object received at the access point 70, the access point 70 may forward the request to the satellite gateway 72. By delaying transmission of the request, the number of requests transmitted over the wireless link may be reduced.

5 In alternative embodiments, these problems of correlating requests from the web browser 20 with the objects which are in the process of being prefetched and forwarded in a chubby response can be avoided. In one embodiment, the access point 70 parses the incoming files to determine the presence of inline objects in a similar manner as the satellite gateway 72. The access point 70 assumes that the satellite gateway 72 will prefetch all of these inline objects. For example, when the parent file is received in
10 block 88 at the access point 70, the access point 70 may parse the parent file to extract any external references to inline objects. The access point 70 may assume that the gateway 72 is in the process of prefetching the inline objects and may create its own list
15 of prefetch objects comprising all of the inline objects. Because the access point 70 always receives the parent file before the web browser 20, the access point 70 may establish its list of prefetch objects before any requests are received from the web browser 20. If a request is received which appears on the list of prefetch objects but the object itself is not received at the access point 70 within a given amount of time, the access point 70 may forward the request to the satellite gateway 72. Notice that in this
20 embodiment, the transfer of notification data from the satellite gateway 72 to the access point 70 noted in block 92 of Figure 7 does not occur, thus, decreasing the amount of data transferred over the satellite link. Because the correlation problem is avoided, the access point 70 need not delay asking for objects should it receive a request for an object which is neither available nor on the list, thus, decreasing the overall latency
25 associated with system use.

In yet another alternative embodiment, the correlation problem is addressed without increasing the functionality of the access point 70 to include parsing. Referring again to Figure 7, when the satellite gateway 72 receives parent file from the Internet 24 in block 86, it may first parse the parent file to determine whether there are any external references to inline objects before forwarding the parent file information to the access point 70. If there are no external references, the satellite gateway 72 may immediately

forward the parent file to the access point 70. If one or more external references are made, the satellite gateway 72 may request and receive data corresponding to one or more of the inline objects before beginning transfer of the parent file to the access point 70. In this way, the inline objects consistently follow directly after the parent file and, thus, are available at the access point 70 when requested by the web browser 20. In one embodiment, the satellite gateway 72 receives all the inline objects before beginning transmission of the parent file to the access point 70. In another embodiment, the satellite gateway 72 receives one or more of the inline objects before beginning transmission of the parent file to the access point 70. In yet another embodiment, the satellite gateway 72 receives a predetermined amount of data which may correspond to a portion of one or more inline objects before beginning transmission of the parent file to the access point 70. In addition to avoiding the correlation problem, this method does not require that notification messages are sent from the satellite gateway 72 to the access point 70, thereby, reducing the traffic flow over the wireless link. Because the correlation problem is avoided, the access point 70 need not delay forwarding requests for objects should it receive a request for an object which is not available, thus, decreasing the overall latency associated with system use.

The immediate availability of the inline objects at the web browser 20 after display of the parent web page in each of these embodiments has the advantage of providing the appearance of instantaneous response to the human user.

In yet another embodiment, upon receipt of a parent file at the satellite gateway 72, the satellite gateway 72 parses the parent file and begins to request inline objects. As the inline objects are received, the satellite gateway 72 forwards the inline objects to the access point 70. After the satellite gateway 72 has requested, received and transferred all or a portion of the inline objects, including the nested inline objects within other inline objects, the satellite gateway 72 begins to forward the parent file to the access point 70. The transfer of the prefetch inline objects before the parent file is transferred assures that all inline objects are available at the access point 70 before they are requested by the web browser 20. In addition, this embodiment does not require the use of notification messages nor the inclusion of additional functionality at the access

point 70 and allows the access point 70 to forward requests for unknown objects without insertion of intentional delay.

As noted above, data is typically transferred over current standard digital data networks typically according to the HTTP protocol. HTTP is a request/response protocol. The client sends a request to the server and the server answers with a response. This simple protocol avoids the use of multi-step handshake routines in the beginning of each data transmission. However, there is no provision in the HTTP protocol for a client to receive an unrequested data transfer in a chubby response message.

In the preferred embodiment, the link which connects the satellite gateway 72 to the access point 70 comprises a means of transferring unrequested data using a chubby response message. To provide such functionality, the link which connects the satellite gateway 72 to the access point 70 is outside of the pure HTTP protocol. It is important that the distributed proxy server 68 be transparent to the web browser 20 and web server 26. For this reason, it is important that the distributed proxy server 68 support well-known data transfer protocols such as HTTP, at least as perceived by external entities. However, because the satellite gateway 72 and the access point 70 are peer units controlled by a common access provider, it is not necessary that the link which connects the access point 70 and the satellite gateway 72 operate according to industry standard protocols. Therefore, one skilled in the art may develop a protocol according to well-known techniques which allows unrequested data transfer to occur between the satellite gateway 72 and the access point 70 such as through the use of a chubby response.

In Figure 6, a single web browser 20 is shown as coupled to the distributed proxy server 68. However, in the preferred embodiment, the access point 70 serves a plurality of users. In addition, the satellite gateway 72 may serve a plurality of access point 70 units. In one embodiment, the access point 70 is consolidated within a single housing. The housing is coupled to an antenna which is positioned external to a building in which the housing is placed. A local area network couples a plurality of web browser 20 to the housing. In another embodiment, the access point 70 implemented on same machine or housing in which the web browser 20 operates. In yet another embodiment, the functionality and modules of the access point 70 are

5

distributed across a plurality of computing units. In general, the functions and modules of the access point may be distributed across entities on one side of the subject link and the gateway functions and modules may be distributed across entities on the other side of the link. It is not necessary that either the access point or the gateway are directly coupled to the subject link and intermediate entities as well as the subject link may separate the access point from the gateway.

10

15

20

The same techniques which are described above can be applied to other external files even if the probability of a subsequent request for those files is less than the probability of requests for the inline objects. For example, many web pages contain hypertext links to other web pages which the user may choose to access. A home page for a manufacturing company may have a hypertext link to an employment page, a product information page, a company information page and a press release page. A user who accesses the home page is reasonably likely to access one of the pages referenced by the hypertext links. After all the inline objects are retrieved by the satellite gateway 72, the gateway may begin predictive prefetching of the linked web pages and transferring them to the access point 70 in a similar manner as the inline objects. The amount of data may be limited so as to prevent an endless cascade of data. If the user requests one of the predictively prefetched links, he receives the information directly from the access point 70, thereby, avoiding the round trip delay associated with using the wireless link. In general, the prefetched objects may be text or digital data files, images, sound clips, program modules, package of data with external functions or any such item.

25

30

Although the example above was described in terms of a web browser accessing web server over the Internet, the principles of the invention may be applied in many situations where a data terminal requests a data file which contains a reference to an external data file over a data network having a link which introduces substantial delay. In such a case, the requesting unit may be any sort of digital equipment which can generate a file request and the responding unit may be any type of digital equipment which can transfer a file. For example, the Internet 24 shown in the figures may be a public or private local area network or wide area network other than the Internet.

A myriad of alternate embodiments will be readily apparent to one skilled in the art upon examination of the disclosure. For example, although the preferred embodiment was shown with reference to a satellite link, the teachings of the invention may be directly applied to any link which introduces significant delay whether a wire
5 line or wireless link.

In addition, as noted above, the use of the invention reduces amount of data which flows from the access point to the gateway. This reduction can be important even in those systems where the delay of the path between the access point and the gateway is not substantial. For example, cable modems are highly asymmetrical links. The
10 down stream link from the Internet to the user has a fairly high bandwidth while the upstream link from the user to the Internet has a very limited bandwidth. Therefore, even if delay is not an issue, it may be advantageous to incorporate a distributed proxy server in order to reduce the upstream traffic flow.

The invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiment is to be considered in all respects only as illustrative and not as restrictive and the scope of the invention is, therefore, indicated by the appended claims rather than the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.
15

WHAT IS CLAIMED IS:

1. In an internet access system that includes a communications link, the communications link having a client side and having a server side which is remote from the client side, a method of processing a client request for a document that includes an inline object, the client request generated by a client and directed to a document server, the document comprising a base component which includes a reference to the inline object, the method comprising:

on the client side of the communications link:

(a) receiving the client request and forwarding the request over the communications link to the server side;

on the server side of the communications link:

(b) receiving the client request, and forwarding the request to the document server over an internet;

(c) receiving the base component from the document server, parsing the base component to identify the reference to the object, and forwarding the base component over the communications link to the client side for delivery to the client; and

(d) prefetching the inline object from the document server using the reference identified in step (c), and without waiting for the client to request the object, forwarding the object over the communications link to the client side for delivery to client.

2. The method of Claim 1, further comprising the steps of, on the client side of the communications link:

receiving and caching the object forwarded in step (d); and

in response to a client request for the object, forwarding the object to the client.

3. The method of Claim 2, further comprising the step of intercepting said client request for the object on the client side of the communications link to prevent the client request for the object from being transmitted over the communications link.

5 4. The method of Claim 1, wherein the communications link is a satellite link, and the method avoids a delay normally associated with sending a client request for the object over the satellite link.

10 5. The method of Claim 4, wherein step (a) is performed by an executable component which runs at a user access point to the satellite link, and steps (b)-(d) are performed by a satellite-to-internet gateway component.

6. The method of Claim 1, further comprising repeating step (d) for each of a plurality of additional inline objects of the document.

15 7. The method of Claim 1, wherein step (a) comprises communicating with a web browser program using a standard internet protocol.

20 8. The method of Claim 1, wherein the document includes a hypertextual link to a second document, and the method further comprises prefetching the second document on the server side and forwarding the second document to the client side.

25 9. In a client-server type document retrieval system in which inline objects of documents are requested and retrieved separately from base components of the documents, a distributed system for reducing a performance degradation caused by a communications link, the distributed system comprising:

30 a first component which runs on a client side of the communications link and communicates with clients, the first component adapted to receive document requests from the clients and to forward the requests over the communications link for processing; and

5 a second component which runs on a server side of the communications link and communicates with document servers, the second component adapted to receive the document requests from the first component over the communications link and to forward the requests to the document servers, the requests causing the document servers to return base components of requested documents;

10 wherein the second component processes base components returned by the document servers by at least (i) parsing the base components to identify references to inline objects, (ii) prefetching the inline objects, and (iii) forwarding the prefetched inline objects to the first component without waiting for client requests for the inline objects;

15 and wherein the first component stores prefetched inline objects received from the second component in a cache memory, and responds to object requests from the clients by forwarding the inline objects to the clients from the cache memory.

20 10. The distributed system of Claim 9, wherein the first component intercepts a client request for an object when the object resides within the cache memory, the first component thereby reducing traffic over the communications link.

25 11. The distributed system of Claim 9, wherein the communications link is a satellite link, and the first and second components reduce a delay normally associated with requests for inline objects over the satellite link.

30 12. The distributed system of Claim 11, wherein the first component runs at a client access point to the satellite link, and second component is part of a satellite-to-internet gateway system.

13. The distributed system of Claim 9, wherein the communications link is a cable-based link for providing internet access by cable modem.

14. The distributed system of Claim 9, wherein the document retrieval system is a hypertextual system, and the second component is further configured to prefetch a document that is hypertextually referenced within a requested document.

5 15. The system of Claim 9, wherein first component and the second component implement standard World Wide Web protocols for communicating with the clients and document servers, respectively.

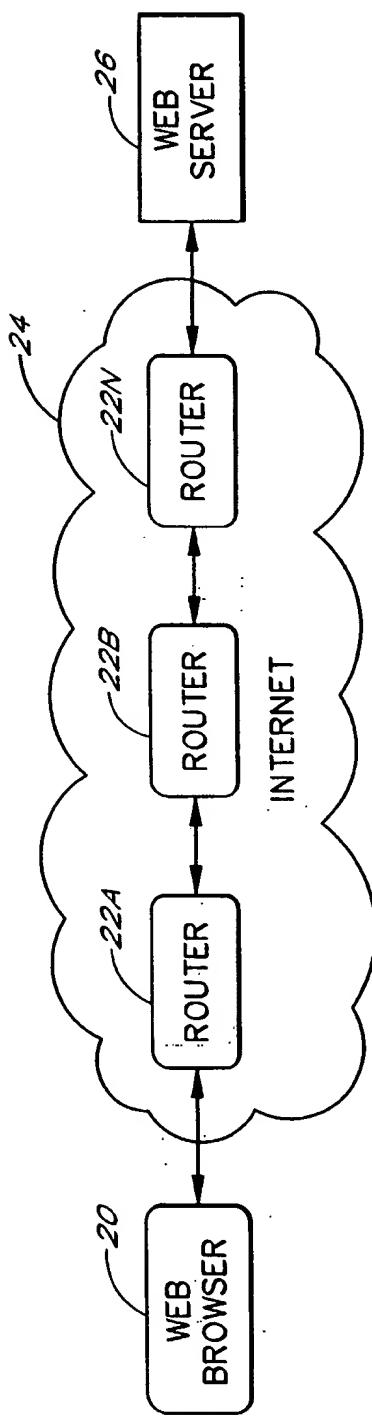


FIG. 1

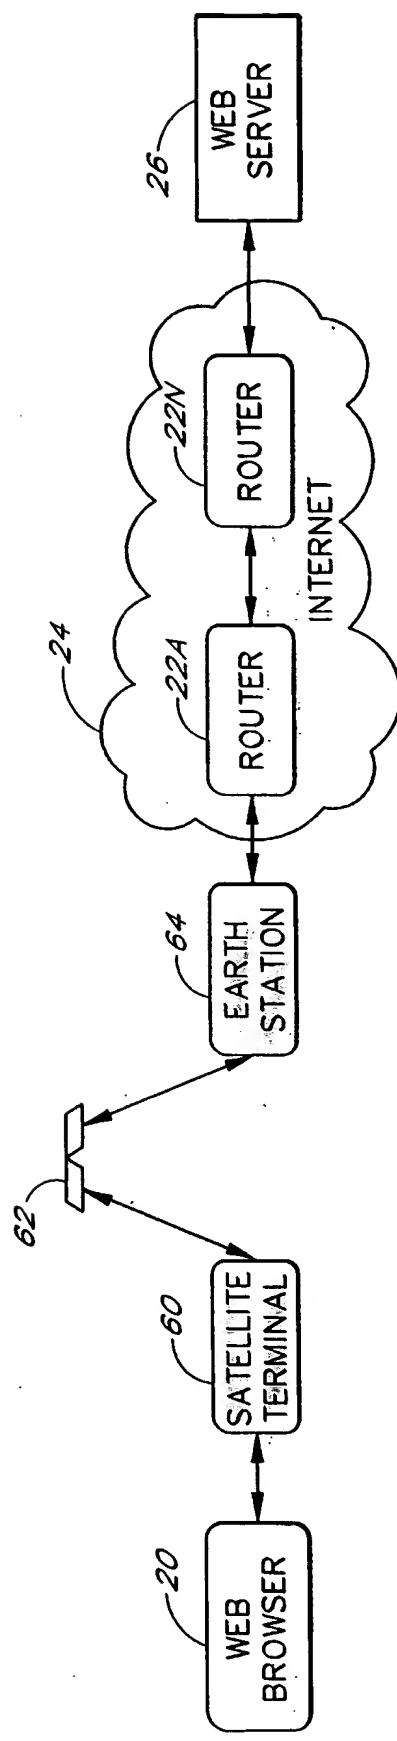
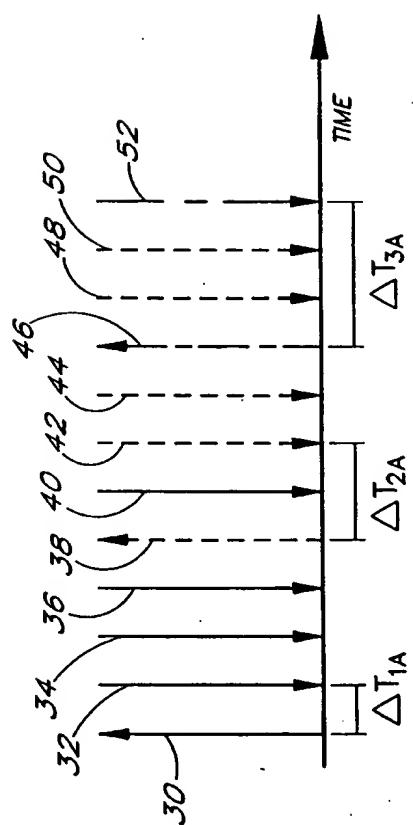
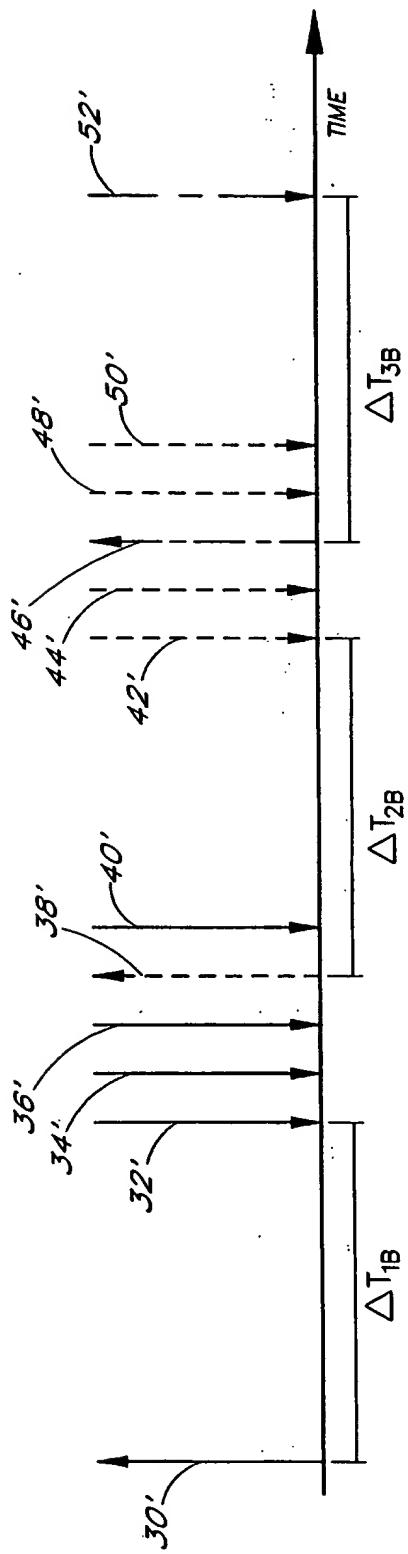


FIG. 3



F/G. 2



F/G. 4

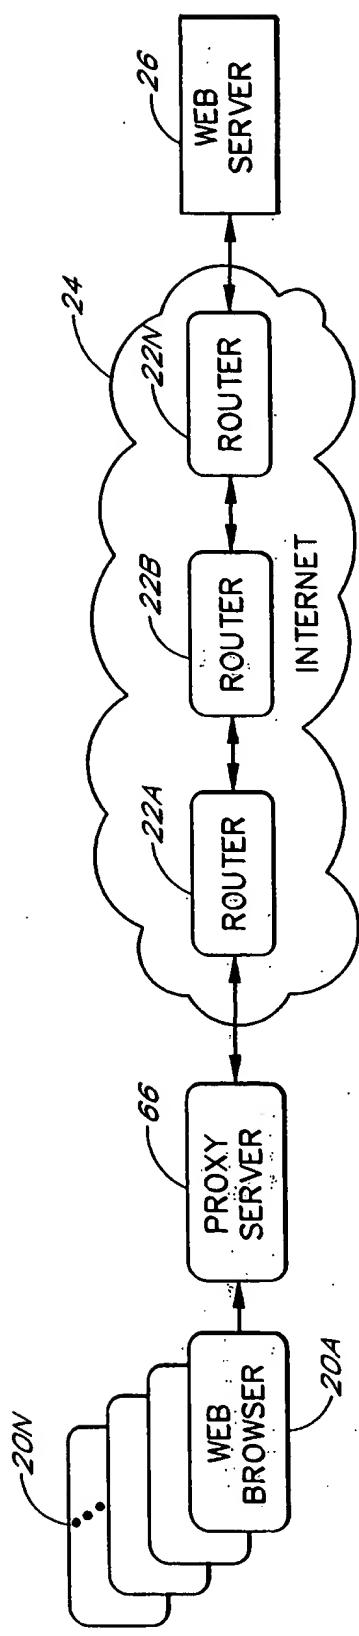


FIG. 5

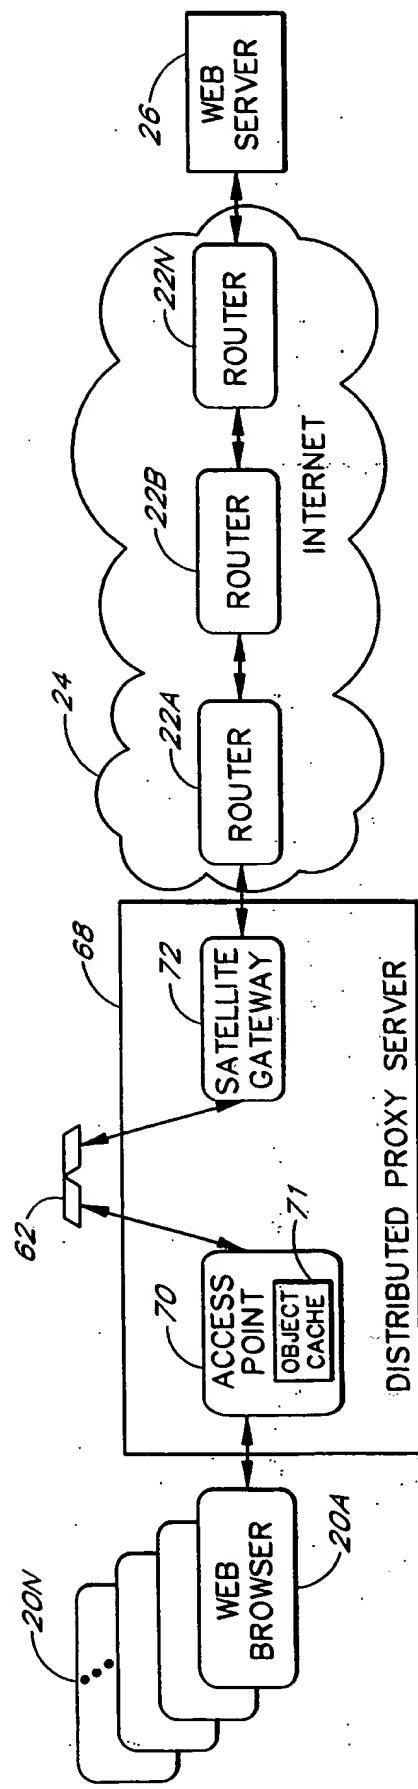


FIG. 6

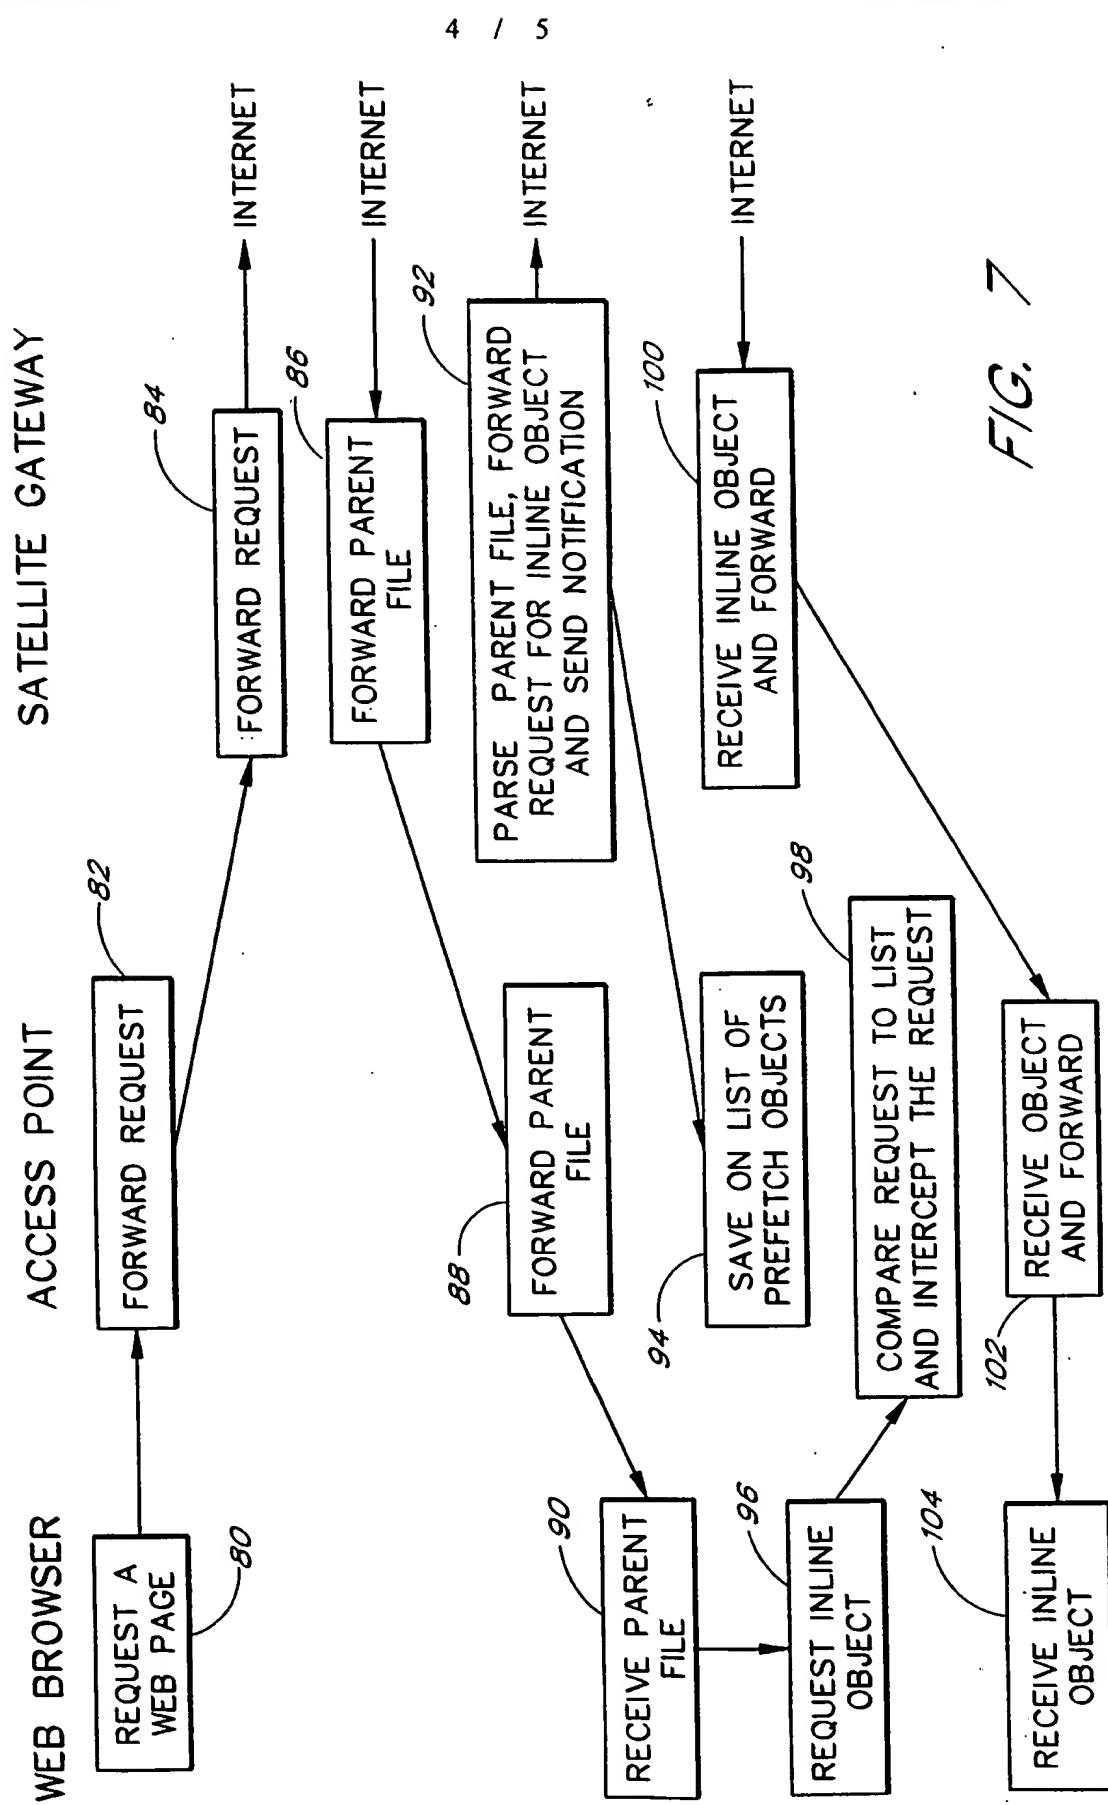
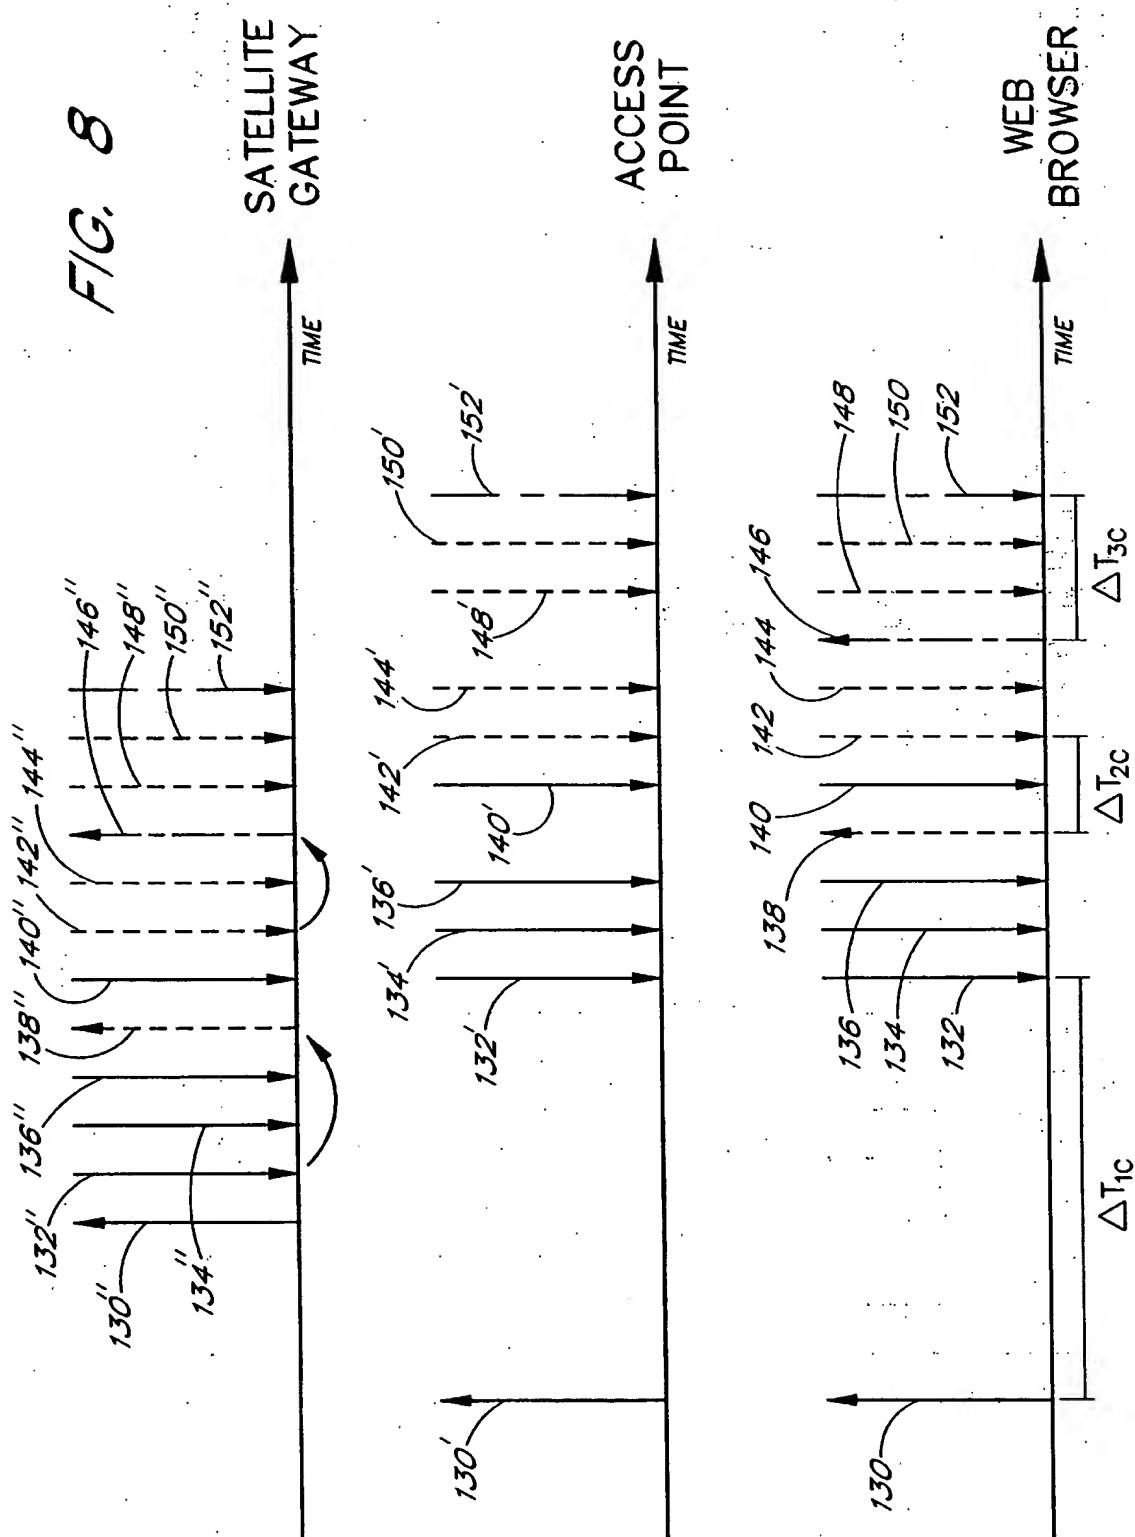


FIG. 7



INTERNATIONAL SEARCH REPORT

Int'l. Jonal Application No
PCT/US 98/16353

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L29/08 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>WANG Z ET AL: "PREFETCHING IN WORLD WIDE WEB" COMMUNICATIONS: THE KEY TO GLOBAL PROSPERITY. GLOBECOM 1996 GLOBAL INTERNET 96 CONFERENCE RECORD, LONDON, NOV. 18 - 22, 1996, vol. SUPP, 18 November 1996, pages 28-32, XP000741668</p> <p>INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS</p> <p>see abstract</p> <p>see page 28, paragraph 1</p> <p>see page 28, paragraph 3</p> <p>see page 30, left-hand column, line 35, paragraph 4 - line 48</p> <p>see page 30, right-hand column, line 30, paragraph 5 - page 31, left-hand column, line 46, paragraph 6</p> <hr/> <p style="text-align: center;">-/-</p>	1-15

Further documents are listed in continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

11 January 1999

Date of mailing of the international search report

20/01/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patenttaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo rd,
Fax: (+31-70) 340-3016

Authorized officer

Adkhis, F

INTERNATIONAL SEARCH REPORT

Inte onal Application No
PCT/US 98/16353

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>SHRIKUMAR H ET AL: "Thinternet: life at the end of a tether" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 27, no. 3, December 1994, page 375-385 XP004037970 see abstract see page 382, left-hand column, line 48 - right-hand column, line 10 see page 383, left-hand column, line 9 - line 21 ---</p>	1-15
X	<p>C. BROOKS, M. S. MAZER, S. MEEKS AND J. MILLER: "Application-specific Proxy Servers as HTTP Stream Transducers" WWW4 CONFERENCES, December 1995, pages 1-9, XP002081461 http://www.w3.org/Conferences/WWW4/Papers/56/ see abstract see page 1, paragraph 1 - page 2, paragraph 2 ---</p>	1-15